CALL FOR PAPERS
3nd International Workshop on
Fault-Tolerance for HPC at Extreme Scale (FTXS 2013)

In conjunction with
The 22nd International ACM Symposium on
High Performance Parallel and Distributed Computing (HPDC 2013)
New York City, New York, USA on June 17-21, 2013 (workshop on June 18th)

WORKSHOP MOTIVATION
For the HPC community, a new scaling in numbers of processing elements has superseded the historical trend of Moore's Law scaling in processor frequencies. This progression from single core to multi-core and many-core will be further complicated by the community's imminent migration from traditional homogeneous architectures to ones that are heterogeneous in nature. As a consequence of these trends, the HPC community is facing rapid increases in the number, variety, and complexity of components, and must thus overcome increases in aggregate fault rates, fault diversity, and complexity of isolating root cause.

Recent analyses demonstrate that HPC systems experience simultaneous (often correlated) failures. In addition, statistical analyses suggest that silent soft errors can not be ignored anymore, because the increase of components, memory size and data paths (including networks) make the probability of silent data corruption (SDC) non-negligible. The HPC community has serious concerns regarding this issue and application users are less confident that they can rely on a correct answer to their computations. Other studies have indicated a growing divergence between failure rates experienced by applications and rates seen by the system hardware and software. At Exascale, some scenarios project failure rates reaching one failure per hour. This conflicts with the current checkpointing approach to fault tolerance that requires up to 30 minutes to restart a parallel execution on the largest systems. Lastly, stabilization periods for the largest systems are already significant, and the possibility that these could increase in length is of great concern. During the Approaching Exascale report at SC11, DOE program managers identified resilience as a black swan - the most difficult under-addressed issue facing HPC.

OPEN QUESTIONS
What does the fault-tolerance community need to do in order to be prepared to face the challenges of extreme scale computing? What is needed to keep applications with billions of threads of parallelism up and running on systems that fail tens of times per day? As models predict less than 50% efficiency of traditional checkpoint/restart methods on future systems, are we ready to pay the cost of full redundancy, effectively performing redundant multi-threading (RMT) across entire systems? Do we even have the infrastructure necessary to implement an RMT strategy?

How is the supercomputing community going to efficiently isolate failures on enormously complex systems? Is there any chance to understand these systems in such a way that some failure could be predicted with enough accuracy and anticipation to trigger useful failure avoidance actions? What can the community do to protect applications from SDC in memory and logic? How far the user and the programmer should be involved in managing faults? What are the most promising self-healing numerical methods?

GOALS
The goals of this workshop are to consider these complex questions, to discuss the unique limitations that extreme scale and complexity impose on traditional methods of fault-tolerance, and to explore new strategies for dealing with those challenges.

PAPER SUBMISSIONS
Submissions are solicited in the following categories:

- Regular papers presenting innovative ideas improving the state of the art.

- Experience papers discussing the issues seen on existing extreme-scale systems, including some form of analysis and evaluation.
- Extended abstracts proposing disruptive ideas in the field, including some form of preliminary results

Submissions shall be sent electronically, must conform to IEEE conference proceedings style and should not exceed eight pages including all text, appendices, and figures.

TOPICS

Assuming hardware and software errors will be inescapable at extreme scale, this workshop will consider aspects of fault tolerance peculiar to extreme scale that include, but are not limited to:

- Quantitative assessments of cost in terms of power, performance, and resource impacts of fault-tolerant techniques, such as checkpoint restart, that are redundant in space, time or information
- Novel fault-tolerance techniques and implementations of emerging hardware and software technologies that guard against silent data corruption (SDC) in memory, logic, and storage and provide end-to-end data integrity for running applications; Studies of hardware / software tradeoffs in error detection, failure prediction, error preemption, and recovery
- Advances in monitoring, analysis, and control of highly complex systems
- Highly scalable fault-tolerant programming models
- Metrics and standards for measuring, improving and enforcing the need for and effectiveness of fault-tolerance
- Failure modeling and scalable methods of reliability, availability, performability and failure prediction for fault-tolerant HPC systems
- Scalable Byzantine fault tolerance and security from single-fault and fail-silent violations
- Benchmarks and experimental environments, including fault-injection and accelerated lifetime testing, for evaluating performance of resilience techniques under stress

IMPORTANT DATES

Submission of papers: February 11th, 2013
Author notification: March 18th, 2013
Camera ready papers: April 15th, 2013
Workshop: June 18th, 2013

WORKSHOP ORGANIZERS

Nathan DeBardeleben - Los Alamos National Laboratory
Jon Stearley - Sandia National Laboratories
Franck Cappello - INRIA & University of Illinois at Urbana Champaign

PROGRAM COMMITTEE

Rob Aulwes – Los Alamos National Laboratory
Aurélien Bouteiller – University of Tennessee, Knoxville
Greg Bronevetsky – Lawrence Livermore National Laboratory
Clayton Chandler – Department of Defense
Robert Clay – Sandia National Laboratories
John Daly – Department of Defense
Christian Engelmann – Oak Ridge National Laboratory
Felix Salfner – SAP Innovation Center Potsdam
Kurt Ferreira – Sandia National Laboratories
Ana Gainaru – University of Illinois at Urbana-Champaign
Leonardo Bautista Gomez – Tokyo Institute of Technology
Hideyuki Jitsumoto – The University of Tokyo
Rakesh Kumar - University of Illinois, Urbana-Champaign
Zhiling Lan – Illinois Institute of Technology
Naoya Maruyama – Tokyo Institute of Technology
Kathryn Mohror – Lawrence Livermore National Laboratory

Bogdan Nicolae – IBM Research – Ireland
Rolf Riesen – IBM Research – Ireland
Yve Robert – ENS Lyon
Thomas Ropars – EPFL
Mitsuhisa Sato – University of Tsukuba
Stephen Scott – Tennessee Tech University and Oak Ridge National Laboratory
Vilas Sridharan – AMD, Inc.
Roel Wuyts – Intel ExaScience Lab

MORE INFORMATION
See http://institute.lanl.gov/resilience/workshops/ftxs/ and http://hpdc.org/2013/ for more information.